

WE CLAIM:

1. A method of allocating bandwidth in a computer network, the method comprising:

determining a first time at which a first token bucket controlling a first bandwidth allocation is approximately full of first tokens; and

allocating, after the first time, excess first tokens to a second token bucket controlling a second bandwidth allocation.

2. The method of claim 1, wherein:

the first token bucket has a first capacity;

the second token bucket has a second capacity; and

the second capacity increases after the first time.

3. The method of claim 1, wherein the first bandwidth allocation and the second bandwidth allocation are dedicated to a single customer.

4. The method of claim 1, wherein the first bandwidth allocation is dedicated to a first customer and the second bandwidth allocation is dedicated to a second customer.

5. The method of claim 1, further comprising:

determining a second time at which the second token bucket is approximately full of second tokens and excess first tokens; and

allocating, after the second time, excess second tokens to a third token bucket controlling a third bandwidth allocation.

6. The method of claim 1, wherein the first bandwidth allocation and the second bandwidth allocation comprise committed information rates.

5 7. The method of claim 1, wherein the first bandwidth allocation and the second bandwidth allocation comprise peak information rates.

8. The method of claim 1, wherein the first token bucket and the second token bucket are color-aware token buckets.

10 9. The method of claim 1, further comprising:

determining a second time at which the second token bucket is approximately full of second tokens and at which the first token bucket is not full of first tokens; and

15 allocating, after the second time, additional second tokens to the first token bucket.

10. The method of claim 9, wherein:

the first token bucket has a first capacity;

the second token bucket has a second capacity; and

20 the first capacity increases after the second time.

11. A method of allocating bandwidth in a computer network, the method comprising:

determining that first tokens added to a first token bucket controlling a first bandwidth allocation will be excess first tokens; and

25 allocating the excess first tokens to a second token bucket controlling a second bandwidth allocation.

12. A method of allocating bandwidth in a computer network, the method comprising:

determining that a first token bucket controlling a first bandwidth allocation is approximately full of first tokens;

determining a state of network congestion; and

allocating, when the state of network congestion is at an acceptable level, additional first tokens to a second token bucket controlling a second bandwidth allocation.

13. The method of claim 12, further comprising the step of allocating, when the state of network congestion is not at an acceptable level, fewer than all additional first tokens to the second token bucket.

15 14. A computer program embodied in a machine-readable medium, the computer program controlling a network device to perform the following steps:

determining a first time at which a first token bucket controlling a first bandwidth allocation is approximately full of first tokens; and

allocating, after the first time, excess first tokens to a second token bucket controlling a second bandwidth allocation.

25 15. The computer program of claim 14, wherein the first bandwidth allocation and the second bandwidth allocation are dedicated to a single customer.

16. The computer program of claim 14, wherein the first bandwidth allocation is dedicated to a first customer and the second bandwidth allocation is dedicated to a second customer.

17. The computer program of claim 14, the computer program further controlling the network device to perform the following steps:

5           determining a second time at which the second token bucket is approximately full of second tokens and excess first tokens; and

allocating, after the second time, excess second tokens to a third token bucket controlling a third bandwidth allocation.

10         18. The computer program of claim 14, wherein the first bandwidth allocation and the second bandwidth allocation comprise committed information rates.

15         19. The computer program of claim 14, wherein the first bandwidth allocation and the second bandwidth allocation comprise peak information rates.

20         20. The computer program of claim 14, wherein the first token bucket and the second token bucket are color-aware token buckets.

21. The computer program of claim 14, the computer program further controlling the network device to perform the following steps:

25           determining a second time at which the second token bucket is approximately full of second tokens and at which the first token bucket is not full of first tokens; and

allocating, after the second time, additional second tokens to the first token bucket.

30         22. A network device for allocating bandwidth in a computer network, the network device configured to perform the following steps:

determining that a first token bucket controlling a first bandwidth allocation is approximately full of first tokens;

determining a state of network congestion; and

5       allocating, when the state of network congestion is at an acceptable level, additional first tokens to a second token bucket controlling a second bandwidth allocation.